Technical Report
1189

Cyber Network Mission Dependencies

A.E. Schulz
M.C. Kotson
J.R. Zipkin

18 September 2015

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

Prepared for the Department of the Air Force under Air Force Contract FA8721-05-C-0002.

This report is based on studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology. This work was sponsored by the Air Force Life Cycle Management Center (AFLCMC/XZCC) under Air Force Contract FA8721-05-C-0002.

Massachusetts Institute of Technology
Lincoln Laboratory

Cyber Network Mission Dependencies

A.E. Schulz
M.C. Kotson
Group 51

J.R. Zipkin
Group 58

Technical Report 1189

18 September 2015

Approved for public release; distribution is unlimited.

Lexington Massachusetts

This page intentionally left blank.

ABSTRACT

Cyber assets are critical to mission success in every arena of the Department of Defense. Because all DoD missions depend on cyber infrastructure, failure to secure network assets and assure the capabilities they enable will pose a fundamental risk to any defense mission. The impact of a cyber attack is not well understood by warfighters or leadership. It is critical that the DoD develop better cognizance of Cyber Network Mission Dependencies (CNMD). This report identifies the major drivers for mapping missions to network assets, introduces existing technologies in the mission-mapping landscape, and proposes directions for future development.

This page intentionally left blank.

ACKNOWLEDGMENTS

The authors wish to thank Kevin Carter, Chris Degni, Jeffrey Gottschalk, David O'Gwynn, Reed Porada, Pierre Trepagnier and David Weller-Fahy for their valuable input and insight.

This page intentionally left blank.

TABLE OF CONTENTS

This page intentionally left blank.

LIST OF FIGURES

This page intentionally left blank.

1. INTRODUCTION

Ongoing and effective protection of the United States against its adversaries requires that the Department of Defense (DoD) be able to guarantee that it can continue accomplishing critical missions, even in the face of degraded or disabled infrastructure. The DoD has recognized that all missions depend on cyber infrastructure. The 2010 Quadrennial Defense Review [1] finds that "A failure by the Department to secure its systems in cyberspace would pose a fundamental risk to our ability to accomplish defense missions today and in the future." The essential reason is that the cyber and human-cognition domains overlap profoundly with each of the traditional DoD metric spaces of Land, Maritime, Air and Space. It is critical to regard the cyber domain not as a separate warfighting space, but rather an integral component of each of the traditional DoD spaces. Figure 1, from a Canadian Armed Forces report [2], captures an appropriate mental model.

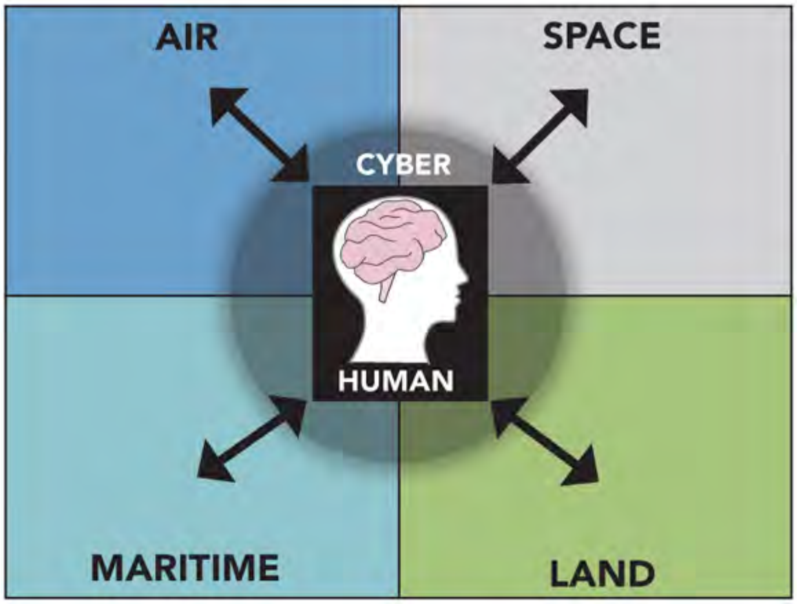

Figure 1. A mental model of the cyber and human-cognition domains in relation to the metric domains of Land, Maritime, Air and Space. This image is from a report of the Canadian Armed Forces [2].

Despite the fact that all DoD missions depend on cyber infrastructure, it was found in the 2008 USAF Science Advisory Board Report on Defending and Operating in a Contested Cyber Domain [3] that "The full range of possible mission effects of cyber attacks is not well understood by warfighters." Ironically, a cyber attack frequently actuates a far greater mission impact than the attacker's direct intent. When faced with evidence of a cyber intrusion or attack, all but the most highly trained defenders will succumb to the human instinct to disconnect the hosts in question, and even turn them off. Not only does this behavior destroy useful evidence for attribution of the attack and discovery of the adversary's intent, it also creates a denial of service for all the missions that depend on these hosts.

It is critical that cyber defenders be provided technology to reveal the mission impact of specific hosts in a contested cyber domain. The SAB report [3] also observed that to "fight through" a contested cyberspace, Major Commands (MAJCOMs) must have prioritized their missions and know the mission dependency on cyber infrastructure. Assets identified as being critical in enabling DoD missions are referred to as the Cyber Key Terrain (CKT). Cyber superiority in a contested cyber environment requires that we

- Prioritize mission and mission-essential functions
- Conduct dependency and risk analysis of network assets
- Engineer robust and resilient architectures for CKT
- Monitor and swiftly detect compromise of CKT

The process of identifying the CKT associated with a particular mission is referred to in this report as Cyber Network Mission Dependencies (CNMD) or more succinctly, mission mapping.

2. MISSION MAPPING AND ITS APPLICATIONS

A useful model of mission mapping is presented in Figure 2. Users and defenders of a network typically have several disjoint types of information available about network usage, ranging from very specific to very broad. These are shown in the first column of Figure 2. At the most specific layer, there are host and network logs and artifacts. More generally, there are a number of network enabled capabilities, such as email, web services, document sharing, etc. At a more general level there are specific processes that network users and defenders execute to do their jobs, and at the most general level, there are the missions that are to be accomplished by the users and defenders.

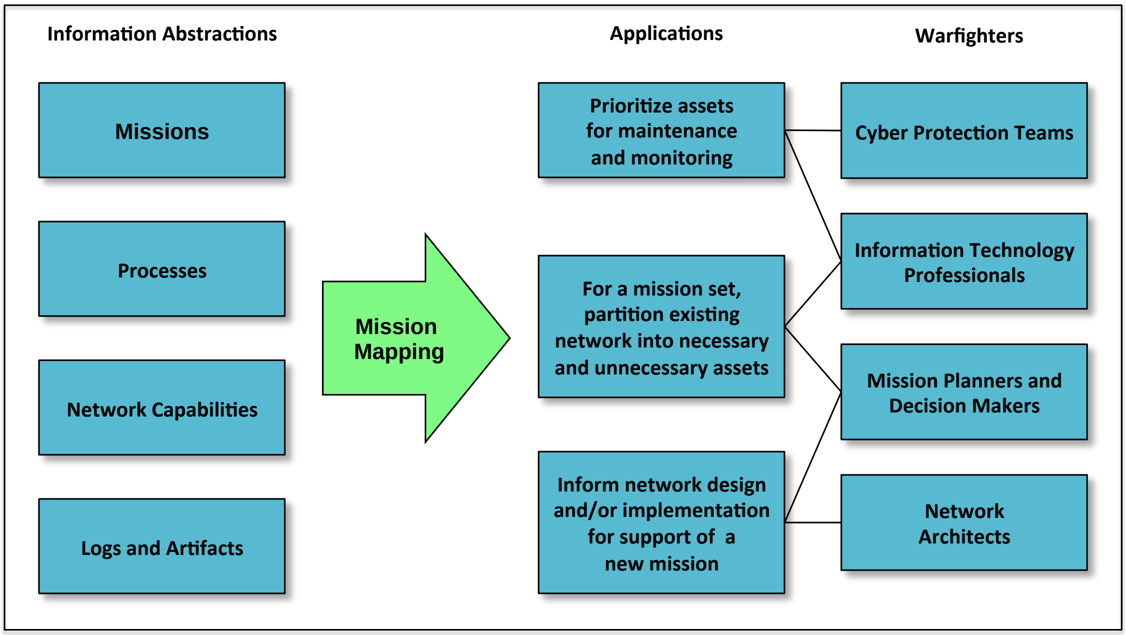
Information Abstractions
Applications
Warfighters
Missions
Processes
Network Capabilities
Logs and Artifacts
Mission Mapping
Prioritize assets for maintenance and monitoring
For a mission set, partition existing network into necessary and unnecessary assets
Inform network design and/or implementation for support of a new mission
Cyber Protection Teams
Information Technology Professionals
Mission Planners and Decision Makers
Network Architects

Figure 2. Information at differing levels of detail is connected together in a dependency map to enable three critical applications: prioritizing assets, partitioning a network, and informing future network design. These applications support Cyber Protection Teams (CPTs), Information Technology (IT) professionals, network architects and decision makers.

Mission mapping supports mission assurance by connecting these layers of information to identify the CKT, the dependencies of the CKT, and the risks associated with the CKT. The logs and artifacts inform us about which capabilities are potentially at risk. Mapping user processes to network capabilities reveals the broader impact of information in the logs, and improves risk analysis by identifying alternative options. For example, a user process may require document sharing. If the network share that is typically leveraged is temporarily degraded in an attack, capability mapping can reveal other options that exist: emailing the document, posting it on a wiki, or serving it up using a web service. The final stage of mission mapping connects the user processes with the missions they support. This mapping is critical both for prioritization of assets given limited time and resources, as well as for understanding what resources will be needed during the planning stages of a mission.

The middle column of Figure 2 summarizes the primary applications that a mission map is used for, and third column describes the job function of the users and defenders these applications support. The first application of a functional mission map will help Cyber Protection Teams (CPTs) and Information Technology (IT) professionals to prioritize assets for maintenance, defense and monitoring. Without this map it is difficult for CPTs and IT professionals to focus attention and limited resources on mission-critical assets. The map will also inform the development of effective policies for the organization.

The second application of the mission mapping tools is to effectively partition the network to separate a particular mission's CKT from other assets that do not directly support that mission. This is particularly useful in a contested cyber environment. Noncritical assets nevertheless increase the available attack surface presented to an adversary. This partitioning will help to quantify the additional risk introduced by assets that are not mission critical. Decision makers may choose to quarantine or disable these tertiary assets to minimize risk of compromise. IT personnel can also use the partitioning to help them with prioritization and resource allocation.

Finally, the third application of mission mapping is to inform network design and implementation for the support of a new mission. Analysis of past requirements provides useful information about asset interdependency, as well as robustness and resiliency. The CNMD mapping can help network architects and mission planners generate cyber requirements during the planning phase of full operations, instead of after the fact. It also helps them quantify risk based on the planned network dependencies. If the risk is too great, these maps can drive investment in new assets or technologies that will be used to help mitigate the risk. It is worth highlighting that a mapping between network capabilities and the supporting assets enables decision makers and mission planners to determine what their network requirements will be without being experts in computers or networking.

3. EXISTING TECHNOLOGIES

The team conducted an extensive investigation into what technologies already exist for CNMD. Some of the technologies we examined were designed specifically for discovery and visualization of mission dependencies. Other tools were not designed for this purpose but provide relevant capabilities that could be leveraged to solve the problem. We researched technologies at MIT Lincoln Laboratory, and also at other Federally Funded Research and Development Centers (FFRDCs), national laboratories, commercial industries, and academic institutions. Specific descriptions of these technologies can be found in Sections 3.2, 3.3 and 3.4. Before discussing specific technologies, we begin in Section 3.1 with a summary of common trends in the technologies we studied. We also provide an overview of technology deficits, necessary capabilities that none of the technologies yet address in a satisfactory way.

3.1 TRENDS AND DEFICITS

The technologies we investigated fell primarily in one of two categories: process-driven analysis and artifact-driven analysis. The process-driven analyses are typically conducted manually by subject matter experts (SMEs), who identify both the mission space and the cyber key terrain that supports the mission. The advantage of this is that the SMEs have a detailed understanding of the missions, the processes, and of the complex system interactions that can occur. The drawback is that the dependencies are discovered manually, typically by interviewing SMEs at different levels of granularity until the actual network assets are reached. This mapping activity is very time consuming and only produces a static map of the cyber key terrain. Another drawback of the process-driven approach is that it relies heavily on appropriately chosen and knowledgeable SMEs. They likely are not able to identify a complete inventory of all of the CKT, and furthermore there will be assets on the network that are not analyzed and do not get added to the map. However, a distinct advantage of this top-down methodology is that since the map is created by humans, it is much easier to interpret and use than maps created with automated processes.

The artifact-driven approaches use logs and data from hosts and network sensors to draw inferences about the usage of network assets. One challenge in this approach is the lack of fidelity in information pertaining to the missions. It is difficult to find appropriate data sources that can reveal the missions, processes, and tasks. The two most prevalent strategies are to link assets to missions by pivoting through the workforce, or to infer the missions through the clustering of assets or employees, coupled with semantic analysis of document or netflow content. The fidelity of the mission determination is highly dependent on the quality of the available data, as well as the quality of the algorithms used to draw the inferences. Maps created with this bottom-up approach can be hard for humans to interpret because the results can be both noisy and incomplete. The advantage is that with suitable algorithms, generating CNMD maps is much faster and less labor intensive. Data-driven techniques enable discovery of new assets as they join the network, or of assets that are repurposed to support a new mission function. They are also able to identify hidden dependencies that may elude even the most savvy SME.

Ultimately, it is likely that neither a process-driven nor an artifact-driven approach will entirely solve the problem. Artifact-driven approaches are dynamic but only show which assets are being used at that point in time. They do not give insight into alternative mechanisms that could also be used to execute the mission. process-driven approaches are manual but more easily reveal alternative mechanisms if certain assets are disabled. Pivoting through capabilities also reduces the knowledge of computer/network specifics required to plan a military operation. A semiautomated combination is likely to be optimal. The flexibility and speed of data-driven methods is required to keep maps of cyber network mission dependencies current and accurate, but human input will always be required to verify the automated solutions, and to add insight as to the true mission impact. One way to combine the approaches may be to use process-driven mapping to generate a mission-to-capability mapping, and a data driven approach to generate a capability-to-network asset mapping. The full CNMD map would be a composite of these two products. Alternatively, a computer-assisted manual process may be the key; for example, mapping by construction by partially automating the mapping of assets to missions as the assets are added to the network. None of the technologies we investigated use composite approaches of this type. Another potential composite approach might be to use automation to fuse expertise from large numbers of people; however, none of the manual approaches we examined appeared to employ any form of crowdsourcing.

In examining technologies developed in industry, academia, FFRDCs, and national laboratories, we identified a few deficits – areas where there is opportunity for better technology development. Interestingly, nearly all the technologies focus on networks that already exist, and there is little attention dedicated to developing better network dependency mapping as new networks are created. It is possible that the network dependency mapping problem could be largely solved with more effective processes or policies, rather than novel technologies. These techniques may not map existing networks, but if applied regularly the network will be mapped within a few years. We also noticed that the existing technologies focused on passive data gathering to establish network dependencies. One of the simplest experiments to determine an asset's impact is to turn it off, but none of the dependency mapping efforts considered actively perturbing an asset's availability to assess its role in network performance.

Another interesting observation is that there are two types of missions that require network mapping: acute missions and chronic missions. Acute missions consist of a focused effort that is coordinated as needs arise, and are usually executed only once. Chronic missions are missions which are executed either continuously or repeatedly. Of the technologies we surveyed, only chronic missions were considered. Few or no resources have been dedicated to identifying the mission dependencies of acute missions. This is a major oversight, because acute missions are often critical to DoD operations, and may require a greater level of mission assurance because in some circumstances, they are executed in reaction to some event or threat to national defense.

In general, the definition of "mission" used by the technologies in our survey tended to be very broad in scope. In some cases, all or nearly all network assets support these broad missions in some way. These broad missions are typically composed of multiple acute missions and chronic submissions or tasks. The smaller subdivisions provide specific capabilities and are limited in time and scope. Sometimes they are not known in advance and occasionally arise on short notice. It would be useful to develop more agile mapping technologies. This would enable researchers to narrow the scope of the dependency analysis to these sub-missions or tasks, many of which align with the network capabilities discussed earlier. The map of a broader mission could then be composed with more fidelity from these smaller capability maps.

Another deficit is that no mapping activity accounts for the time evolution of the network dependency mapping. It is commonly recognized that networks evolve, but less well understood that the missions and mission needs may not be static. Even chronic missions evolve over time; it is dangerous to assume a steady state once a dependency map has been made. Effectively measuring and modeling the time evolution of mission dependencies requires infrastructure to preserve memory of past mission dependencies, and a mapping process that is agile enough to perform a new analysis on time scales shorter than the timescales for the mission or network to evolve.

Finally, none of the technologies we examined have considered the security implications of the network dependency map. There are two primary issues to be considered. First, assuring and reporting the provenance of the data is important for the integrity of the map. This may play a role in how to instrument the network with sensors, and anti-tamper mechanisms will need to be implemented for any reliable CNMD pipeline. The second factor is that once the mission map is created, it will contain all the information needed to cripple our efforts. This intelligence is very dangerous to our missions and is likely to become a target for enemy intelligence operations. Our survey of existing technologies did not identify any solutions to mitigate this problem.

3.2 MIT LINCOLN LABORATORY

MIT Lincoln Laboratory (MIT LL) has recognized the critical need for CNMD and has been active in this problem space for some time. We summarize several of the most noteworthy contributions to this research area.

One of the earliest effort developed at MIT LL was a passive network mapping tool known as NetGlean. In this technology, a wide array of protocols are studied in order to determine the network structure. The protocols include DHCP, DNS, HTTP, various printer identification protocols, and various mail server protocols. Observing this network traffic passively allows NetGlean to discover a complete list of network assets in near-real time without consuming large amounts of bandwidth and interrupting network services. Once a map is completed, NetGlean provides configurable visualizations of the network. No mission context is analyzed, which limits NetGlean's capabilities as a pure mission-mapping tool.

The Network Security Planning Architecture (NetSPA) was developed by MIT LL to automatically and efficiently analyze the security of a network against cyber attacks [4]. This program simulates the activities of a red team or adversarial entity using attack graphs, revealing how nodes on a network are linked together and how these links can be exploited. NetSPA can scan very-large-scale networks (thousands of hosts) in under a minute. The NetSPA Graphical User Interface (GUI) provides illustrations of the network structure and a list of recommended actions that can improve network security. However, NetSPA provides no automated means of identifying network components as important to a given mission.

MIT Lincoln Laboratory has also developed Probabilistic Threat Propagation (PTP) [5], a tool for detecting network nodes which are malicious or have been infected by some malicious entity. PTP uses a graph analysis of the network to determine the probability that each node is a threat, and can identify and characterize groups of malicious nodes (such as a botnet). Using collected email metadata, the same methods can be used to identify the communities to which a person of interest belongs and, by inference, the missions supported by the individual.

AMMO (Automated Mission Mapping Operations) is another mission mapping tool in development. The goal of AMMO is to determine and prioritize the vulnerabilities of cyber assets. It requires a subject matter expert to provide an initial estimate of the cyber key terrain. Mapping is performed by tracking packet data associated with the initial CKT, and inferring an extended list of critical assets based on communications patterns and software dependencies. Once vulnerabilities have been assessed, AMMO produces a prioritized list of patches to be made in order to better protect the network from attack. The data from the risk assessment is also output in a format that can be read by visualization tools such as Dagger (discussed in further detail in Section 3.3).

Person-Centric Automated Mission Mapping, or PCAMM, is a tool developed at Lincoln Laboratory to determine mission-critical information assets [6]. PCAMM first determines the mission structure of an organization using data pulled from directories and financial records, which can identify the programs and projects for which individual employees work. This mission map is then enriched with network data, such as authentication logs, which can connect employees to the information assets which support their work. PCAMM includes many tools for sorting, analyzing, and visualizing these enriched mission map data.

3.3 OTHER FFRDCS AND NATIONAL LABORATORIES

Other FFRDCs and National Laboratories have also been interested in the CNMD problem space. Here we review contributions from MITRE, Lawrence Livermore National Laboratory (LLNL) and Johns Hopkins University Applied Physics Laboratory (JHU/APL).

Cyber Mission Impact Assessment (CMIA) techniques have been developed by MITRE and applied to mission-mapping problems with considerable success [7]. Using modified business process modeling tools, CMIA is able to simulate the components of a mission as a series of tasks with specified durations. The network components required by each task must be determined manually. Once a mission workflow is constructed, CMIA can simulate attacks of various effects and durations to determine how each class of attack affects the integrity of each network machine. In turn, this information can be used to determine how important each machine's availability is to the mission. These tools do not automatically map the network or gather mission details, so they are not ideal for frequent, recurring use to support an organization. Instead, MITRE recommends CMIA be used when planning and prototyping a network, so that the mission impact of various attacks can be understood and minimized from the start.

Also developed by MITRE, the Risk-to-Mission Assessment Process (RiskMAP) is used to determine critical and vulnerable information assets on a network [8]. RiskMAP models a mission as a hierarchy of tools and goals: network nodes support information assets, which in turn perform tasks in support of overarching business objectives. By modeling the effects of a cyber attack on a network node, RiskMAP can determine whether or not the attack jeopardizes an organization's mission. While most network mapping tools can determine the integrity and availability of a cyber asset, RiskMAP is also able to assess the confidentiality of assets, a vital feature to understand for missions geared towards cyber security. Like CMIA, RiskMAP is not designed to automatically map a network or discover new assets, so it is best used as a tool to test prototype networks.

Johns Hopkins University Applied Physics Laboratory has developed Dagger [9], a framework which can model and visualize cyber situational awareness data to reveal mission impacts. To understand the mission and network structure of an organization, Dagger first requires manual input from subject matter experts. Once the mission model is constructed, Dagger is able to pull information from various real-time data feeds and visualize the status of not just network machines, but also software tools, network connections, server room conditions, and many other mission parameters. From this visualization, a user can easily and efficiently observe many layers of abstraction to determine the status of the mission and inform any decisions he or she must make.

Network Mapping System (NeMS) is a software-based tool created by the Lawrence Livermore National Laboratory to discover and map network assets in support of cyber situational awareness [10]. NeMS combines both active probes and passive monitoring of network data to map the network without user input. The mapping process can begin from any node on the network, or from multiple nodes, and the user can adjust the speed, load, and security settings in order to maximize efficiency without disrupting network activities. Tests of NeMS in control networks yielded great results, as the tool identified 100% of known hosts plus an additional previously unknown external connection. Once the mapping is complete, NeMS can also display its own visualization of the network topology. Because it is exclusively a network-mappping utility, NeMS provides no mission context for any discovered network assets.

3.4 INDUSTRY AND ACADEMIA

There are also CNMD contributions of note from industry and academia. The first is CAMUS [11]. Developed by Applied Visions, Inc., CAMUS (Cyber Assets to Missions and Users) is an ontology-based tool which supports many different types of input data. Combining information from several network data sources (including but not limited to FTP logs, Unix logs, and DNS dumps), Camus is able to automatically map the relationships between network machines and the missions they support. It is similar to PCAMM in that it also pivots through the workforce to do so. Camus is even able to detect changes in network or mission structure over time. However, this tool does not provide any information concerning how vital a resource is to mission success or how vulnerable it may be to attacks.

Another interesting technology is NSDMiner. NSDMiner (a loose acronym for "Mining for Network Service Dependencies") is a network-mapping utility developed by North Carolina State University [12]. This tool is meant to automatically discover local-remote dependencies on a network, and it can achieve this without any prior input concerning the network structure. NSDMiner builds its maps using data from network activity, such as Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) flows, but this input data must first be gathered by another service. This utility does not identify information assets in any sort of mission context and is thus poorly suited for automatic mission mapping.

Finally, there is IPsonar [13]. IPsonar is an active network discovery resource produced by Lumeta Corporation. This program is able to scan the networks of globally distributed organizations, evaluating risks on every information asset without consuming enough bandwidth to disrupt business operations. In addition to discovering unknown hosts, IPsonar is able to identify unauthorized connections ("leak paths") and determine if firewalls and router access control lists are violating network policy. Visualization tools are provided to help analysts better understand the network map at multiple levels of detail. No mission mapping utilities are provided.

3.5 COMPARATIVE ANALYSIS

It is useful to assess the relative strengths of the technologies outlined in the previous sections. In Figure 3, we compare the capabilities of the network scanning and mission mapping technologies we reviewed. Each row is a different tool, and each column is a desired mapping functionality. A grey block indicates that a tool is unable to perform the given function. Yellow, green, and blue show the extent to which the function is automated. Yellow processes are dependent on manual data entry from the user, green processes may prompt for user input or provide information that requires further user processing, and blue processes run entirely automatically.

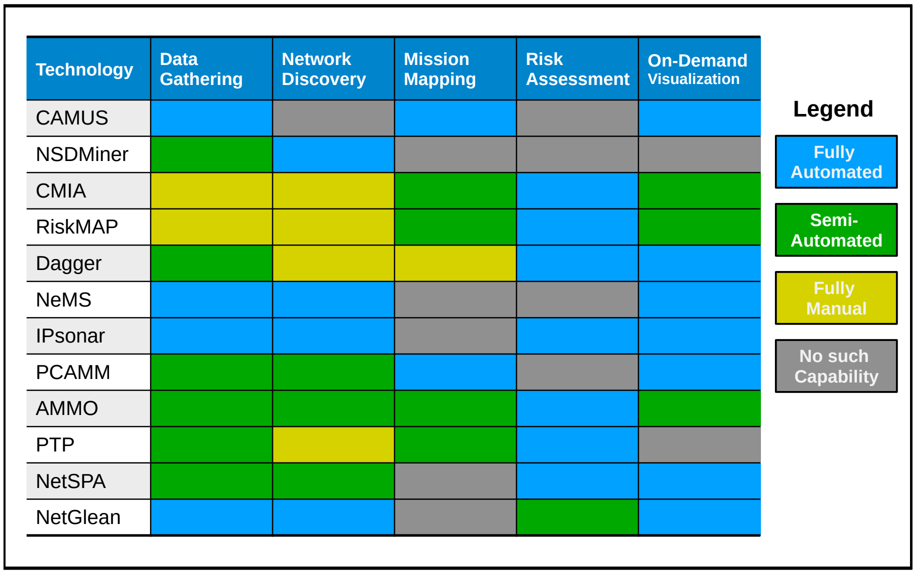

Technology	Data Gathering	Network Discovery	Mission Mapping	Risk Assessment	On-Demand Visualization
CAMUS	🟦	⬜	🟦	⬜	🟦
NSDMiner	🟩	🟦	⬜	⬜	⬜
CMIA	🟨	🟨	🟩	🟦	🟩
RiskMAP	🟨	🟨	🟩	🟦	🟩
Dagger	🟩	🟨	🟨	🟦	🟦
NeMS	🟦	🟦	⬜	⬜	🟦
IPsonar	🟦	🟦	⬜	🟦	🟦
PCAMM	🟩	🟩	🟦	⬜	🟦
AMMO	🟩	🟩	🟩	🟦	🟩
PTP	🟩	🟨	🟩	🟦	⬜
NetSPA	🟩	🟩	⬜	🟦	🟦
NetGlean	🟦	🟦	⬜	🟩	🟦

Legend

🟦	Fully Automated
🟩	Semi-Automated
🟨	Fully Manual
⬜	No such Capability

Figure 3. A comparison chart of the capabilities demonstrated by the technologies we have reviewed. Each row represents a mission mapping or network scanning tool, and each column is a function we expect an ideal mapping tool to perform. The color of each block indicates the extent to which the function is automated.

The capabilities included in the comparison chart are those that we considered most important for a mapping tool. “Data Gathering” refers to the software’s overall ability to obtain the information used for all of its functions. Most tools can obtain at least some of this data automatically, usually by scanning network activity. CMIA and RiskMAP require users to manually input network and mission data, which makes both technologies better suited for mission prototyping than for tracking current mission status. “Network Discovery” and “Mission Mapping” explain the tools’ abilities to infer the structure of an organization’s computer network and to connect the network’s information nodes to mission programs, respectively. Several tools in the chart were designed exclusively for network mapping, and only CAMUS and PCAMM are able to automatically map network assets to missions. “Risk Assessment” is the ability to judge the vulnerability of network nodes to various forms of cyber attack or exploitation. Finally, “On-Demand Visualization” explains whether or not the tool provides a means of viewing the results of its mapping. Some tools, like AMMO, yield data outputs which can easily be displayed with other visualization tools such as Dagger.

In Figure 4 we compare how well each of the technologies we reviewed aligns to the primary mission mapping applications discussed in Section 2. Again, each column represents a mission mapping or network scanning tool, and each row is one of the three major applications expected of a mission mapping technology: the ability to rank network assets in order of their need for monitoring and upkeep, the ability to divide the network into those assets necessary and unnecessary for a given mission, and the ability to inform and guide the user’s decisions when planning a new mission or designing a new network.

It is readily apparent from this chart that no single technology is able to completely satisfy our defined mission mapping needs. Many tools are able to monitor network assets, gauge their importance to the mission, or assess their vulnerability. However, there is no overlap between tools which can completely map a network’s structure and tools which can identify the connection between nodes and mission, so none of these technologies can completely partition the network based on mission necessity. Lastly, few tools seemed to provide a means of designing new networks or missions from the ground up. Notable exceptions are CMIA and RiskMAP, which have been developed by MITRE with network prototyping in mind.

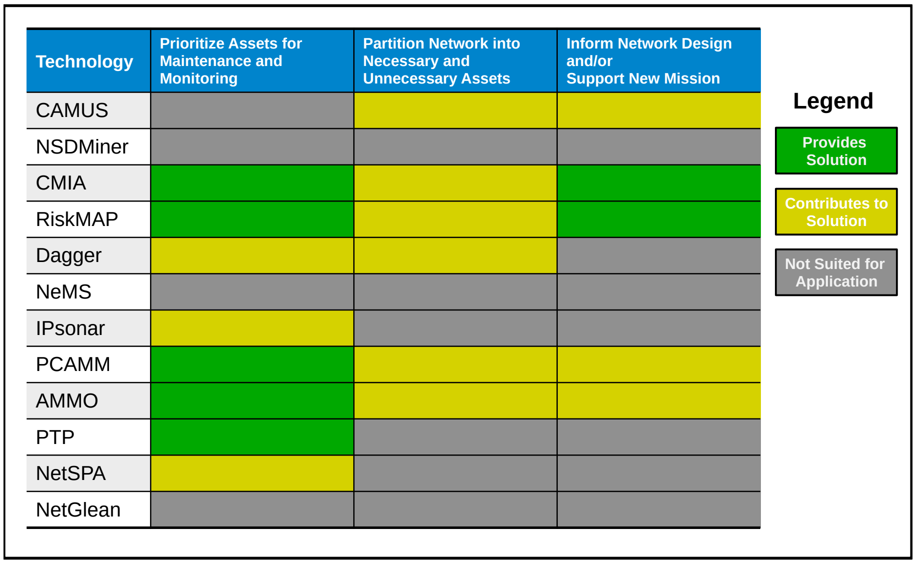

Technology	Prioritize Assets for Maintenance and Monitoring	Partition Network into Necessary and Unnecessary Assets	Inform Network Design and/or Support New Mission
CAMUS	▨	🟨	🟨
NSDMiner	▨	▨	▨
CMIA	🟩	🟨	🟩
RiskMAP	🟩	🟨	🟩
Dagger	🟨	🟨	▨
NeMS	▨	▨	▨
IPsonar	🟨	▨	▨
PCAMM	🟩	🟨	🟨
AMMO	🟩	🟨	🟨
PTP	🟩	▨	▨
NetSPA	🟨	▨	▨
NetGlean	▨	▨	▨

Legend

🟩 Provides Solution

🟨 Contributes to Solution

▨ Not Suited for Application

Figure 4. In this chart, each row represents a mission mapping or network scanning tool, and each column represents one of the three major applications expected of a mission mapping technology. The color code explains whether the tool completely lacks the given application (gray), contributes at least some information relevant to the application (yellow), or provides a complete solution to the application (green).

4. FUTURE TECHNOLOGY DEVELOPMENT

In the following section, we will outline several interesting directions for future research which were not present in our survey of existing technologies.

4.1 MISSION AWARE ARCHITECTURES

It may turn out not to be possible to determine the mission dependencies with high fidelity on an existing network. A promising alternative is to architect network components in such a way as to reveal the mission dependencies. The hope is to partially automate mission mapping as new assets join a network. Such a process would slowly map out missions as network assets are replaced. In a short time ($\sim 5 - 10$ years) the missions would be mapped by construction.

One possible avenue to achieve this end is to use virtualization and software switching to construct a mission-aware network. Virtual Local Area Networks, or VLANs, are used today to achieve defense in depth. The idea is to separate the functional divisions in an organization into separate VLANs, so that if, e.g., the marketing division is compromised, the adversary cannot easily pivot from those hosts to hosts used for other functions such as engineering or accounting. The VLANs could be refined to encapsulate specific network capabilities or missions.

In the past, one VLAN per network capability was impractical because all hosts were physical hosts and all routing occured in hardware. Hosts were multipurposed and contributed to multiple capabilities, and indeed often to multiple mission areas. The advent of mores sophisticated virtualization and software-defined routing mitigates these former obstacles. Virtual Machines (VMs) on the same physical host can exist in different VLANs. Also, software defined routing allows for effortless VLAN redefinition. Therefore, any physical host supporting more than one network capability could do so on separate VMs belonging to separate VLANs. Figure 5 demonstrates this notion visually.

The opportunity for technology development is to research a process for building a mission-aware architecture on a network. Every functioning host on this network will be a single-purpose VM. Every VM supporting a given capability will be assigned to the same VLAN. The VLANs will map to capabilities which will be further grouped to form missions. One aspect of the research will be a trade study. Higher granularity requires dedicating more resources to VM overhead, so it will be necessary to determine the most useful level of granularity. Very fine granularity will cause the number of VMs on the network to burgeon to scales that have not currently been used. Some technology development will need to center around infrastructure to manage and network such a large fleet of VMs, in order to implement the mission-aware architecture with minimal burden on the systems administrators and other IT personnel.

There are interesting security implications of this alternate mission mapping scheme. If unused capabilities and resources are disabled, single-purpose hosts are significantly more defensible than generalized computing platforms, because if compromised, an adversary is limited in their options in using that host for alternate purposes. Furthermore, this network architecture automatically generates tremendous defense in depth.

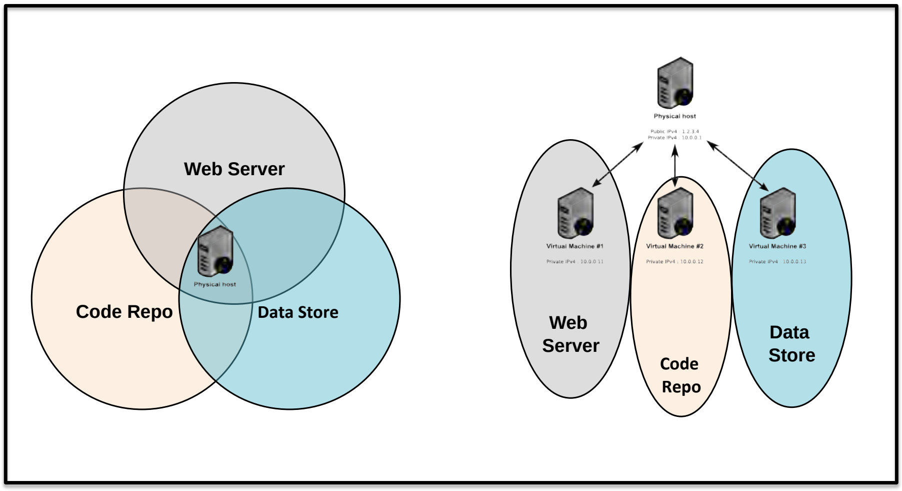

Figure 5. Historically, the same physical host was frequently used to support multiple network capabilities, often in support of multiple missions. A mission-aware architecture could leverage virtualization technology to implement single-purpose VMs that support only one network capability. Virtual hosts associated with the same capability could be grouped together in their own VLAN.

4.2 YOU MAY ALSO LIKE

Missions often require the simultaneous requisition of cyber resources. For example, establishing a deployed email network requires a diverse set of hardware and software that must all be present together at a designated place and time. In addition, some theaters present specific, even unique, cyber threats that demand additional requisitions with which a supply agent may not be familiar. In this environment, errors in requisition are easy to make, and they are costly: an incomplete cyber system may be inoperable, ineffective, or vulnerable while the operators wait for the missing resources. This is a mission mapping problem. Supply agents map the capabilities required by a mission onto the resources needed to provide those capabilities. An error in the mapping can result in an error in the requisition. To decrease the chance of error, supply agents can benefit from access to the practices of supply agents working on the same problem and to any specific requirements their mission faces, as defined by experts elsewhere in the military.

We propose a solution inspired by the “You May Also Like” (YMAL) features of popular services like Amazon and Netflix. A mockup interface is shown in Figure 6. A supply agent would enter the mission details into YMAL. This could be a list of capabilities needed by the mission, or if that space is too large then just a list of already requisitioned resources. YMAL would then apply machine learning techniques to the big data space of past mission requisitions to recommend other resources the mission may require. For example, YMAL could infer that a supply agent ordering a laptop, Ethernet cables, a server, and an Outlook license is establishing an email network and recommend a firewall and additional laptops. YMAL would also match mission details like the deployment location with recommendations from experts.

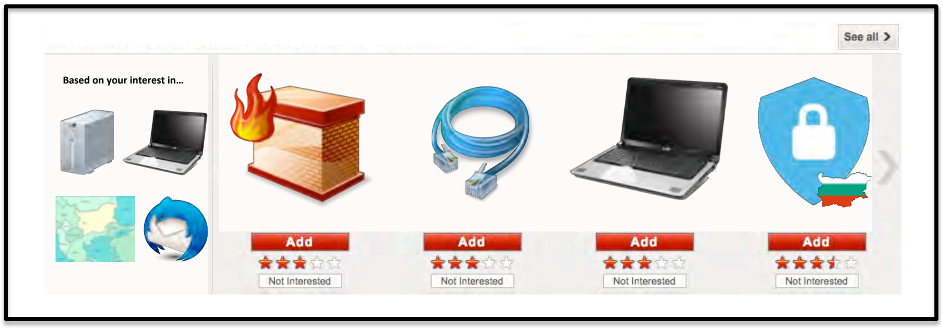

Figure 6. A mockup of an asset recommendation system. Machine learning can be used to make suggestions for potential mission requirements, based on network capabilities required by similar missions in the past.

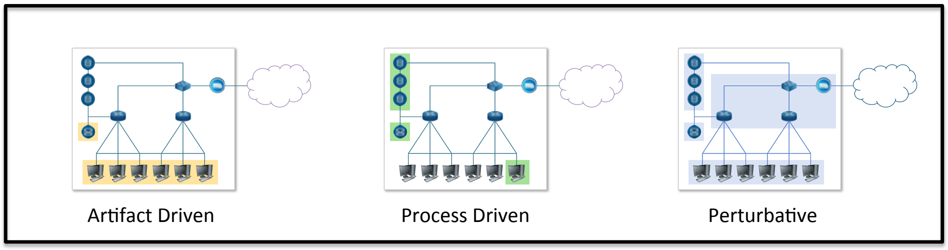

Figure 7. Artifact-driven mapping and process-driven mapping are likely to produce less complete dependency mapping results than an active, perturbative approach.

All stakeholders benefit from this framework. Supply agents get automated assistance choosing resources and staying up to date with common practice and expert recommendations. Commanders face a lower risk of mission impact from requisition errors. Analysts get a rich and expanding source of data for other mission-mapping efforts. Notably, very little of this is specific to cyber. A YMAL system could assist with non-cyber requisitions as well, though the military's institutional familiarity with those domains relative to cyber means that it may provide the most benefit for cyber requisitions. We are currently investigating whether the military has similar systems in place and their status. The primary challenge we anticipate at this point is access control. Allowing data from classified missions into the training set would improve the power of predictions, but it would require ensuring that uncleared personnel have access only to the outputs of the predictions rather than the underlying data. Safeguards against adversarial access not only to the underlying data but also to the higher-level within the recommendation system would be essential.

4.3 ROLE-BASED MISSION BEHAVIOR BASELINE

It is frequently difficult to glean mission context from available network artifacts, which drives the need for manual processes in a functional mission dependency analysis. However, to the extent possible, it is desirable to minimize manual processes because they are labor intensive, and the results of manual analyses can become outdated very quickly. In the DoD, there exist data-driven sources of mission context that have not been leveraged in any of the research we have encountered. Joint mission planning systems such as the Joint Operation Planning and Execution System (JOPES) and the Joint Mission Planning System (JMPS), and similar mission planning systems in the individual services, such as the Air Operations Center (AOC), are a veritable trove of mission-related data that could be leveraged in the mapping between missions, processes and network capabilities. For organizations that employ planning systems, the mapping between network assets and mission impact may be significantly easier to automate.

There is a drawback to leveraging the data available in mission planning systems such as JMPS or the AOC. Very frequently, these data artifacts are classified, which may significantly restrict the set of users who have access to the dependency map. In particular, some IT professionals who support the organization may be unable to leverage the mission map when allocating resources and prioritizing assets. It is desirable to seek a way to construct a mission dependency map without direct access or reference to the mission planning data.

It may be possible to find another way to leverage this data. Individual team members are assigned tasks based on mission planning systems like JOPES/JMPS/AOC. When people execute the task, they reveal which assets support the mission. Mission context can potentially be gleaned by pivoting from tasks through humans to assets leveraged. Assets supporting a mission can be gleaned by baselining the activities of individuals executing tasks in mission planning systems. These signatures can be based on the role the user plays in the mission planning system. This technology is a natural extension to person-centric approaches such as the PCAMM tools at MIT Lincoln Laboratory [6] or CAMUS from Applied Visions [11]. These tools could be adjusted to track metrics on user behavior to assess mission requirements.

The research consists of designing appropriate role definitions for users executing tasks in a mission planning system. Baseline behaviors as a function of user role can be established through an initial learning phase where data on asset usage is recorded. The aggregate behaviors could be used to define role-based signatures for asset usage. The role-based signatures should be observed over a long period of time to quantify whether they are stable, cyclic, or evolve unpredictably with time. If the role-based signatures are stable or cyclic, these can be used to map out the cyber network mission dependencies. The user baselines would also be helpful to detect anomalies in mission execution patterns. Establishing baselines could also be extremely useful when new assets are added to the network, because leveraging changes in the baseline behavior may reveal the new asset's role in the network and missions.

4.4 PERTURBATIVE MISSION MAPPING

The final technology thrust we recommend takes an active approach to generating computer network mission dependency data. This is similar to the approach studied in [14]. The key insight is that if you change the way assets on the network respond to input, you may be able to conduct a sensitivity analysis to reveal how critically a given mission depends on those assets. It may be possible to disrupt or degrade assets on the network, and measure the response, during regular business operations without causing serious harm to the execution of the mission.

This idea is based on the technology of the Simian Army, implemented by Netflix and used very successfully both there and at Amazon [15]. The Netflix Simian Army is designed to force developers to create resilient and robust software architectures that can withstand outages, disruptions, and other chaos that occur on real networks. The Chaos Monkey, a tool that randomly disables production assets, was originally developed to ensure Netflix could continue to deliver pixels to a screen even in the face of common types of failure modes. Other monkeys, such as Latency Monkey and Conformity Monkey, quickly followed once it was realized that live stress testing produced very significant improvements in the performance of the system.

Figure 7 illustrates why perturbative mapping is likely to produce a more complete analysis than either the artifact-driven or the process-driven approaches we have discussed. Suppose authentication logs provide the artifacts needed for a data-driven map. The left panel of Figure 7 represents the terrain that would appear in the resulting CNMD map, highlighted in yellow. The artifacts could probe usage of users' desktop or laptop computers, and potentially mail servers and some database infrastructure if authentication is required. But other parts of the key terrain such as switches or firewalls would not appear in the map because users generally do not directly authenticate to these assets, and no artifacts are generated. The center panel of Figure 7 represents in green the terrain that might be identified when interviewing a subject matter expert about the process of her work. She might correctly identify her workstation as well as some key data stores used in the process, but she may miss routing infrastructure in her critical path, or perhaps infrastructure associated with a "single sign-on" mechanism.

Because the perturbative approach is used to probe mission sensitivity to every asset on the network, the resulting map is significantly more complete than the maps identified with artifact- or process-driven analyses. This is conceptualized in the right panel of Figure 7 in blue. Furthermore, the results of a perturbative analysis are significantly more quantitative, since the sensitivity could be measured as a function of the severity of a disruption. The perturbative approach, if repeated frequently, could also yield insight into the complex system interactions between assets supporting the mission, as well as the impact of other unrelated traffic. Finally, the perturbative approach could scale down to the level of network capabilities much more easily than the other approaches, which would generate better fidelity in the results.

The opportunity for research consists of developing an agent that will introduce artificial latencies or disruptions during the execution of mission-essential functions. It will be particularly interesting to learn whether perturbative disruptions that do not introduce noticeable deficiencies in the operational system can nevertheless be used to probe the sensitivity of user processes to various assets on the network. Developing a mechanism for obtaining feedback on the mission impact will be critical to the success of this technology. Metrics must be devised to quantify the sensitivity of mission execution to asset perturbations. If the impact on the mission execution is non-negligible, it would be wise to conduct a trade-off analysis to balance the insights gained against the temporary inconvenience to the users of the testing process.

5. CONCLUSIONS

Determining the mission impact of a cyber attack on a network, or degradation of service, is one of the highest priorities in assuring the success of DoD missions in defense of the United States. Many organizations are investing in technology development to solve the problem of determining cyber network mission dependencies, and there has been considerable progress. These organizations tend to employ either a manual process-driven approach, or an automated data-driven approach in their technologies. It is our finding that neither fully automated nor fully manual processes appear to be completely effective, and it may be necessary to utilize a combined approach.

All of the technologies we reviewed take a passive data-gathering approach versus an active experimental approach to solving the CNMD problem. We also did not encounter many researchers who attempted to achieve network dependency mapping by construction, as assets are initially added to the network; all the technologies assumed the network in question is already in place and dependencies must be discovered rather than created. We recommend investment in technologies that explore either an active or a constructive approach to mission mapping.

We also observe that current CNMD tools appear to be focusing on the wrong scope; the definition of the mission is often too broad to result in a meaningful mapping to network assets. We advocate shifting the focus to the mapping of network capabilities, and building the broader mission dependency map in reference to these capabilities. Pivoting through capabilities also reduces the knowledge of computer/network specifics required to plan a military operation.

Finally, both networks and missions evolve with time, but few if any of the technologies in development are equipped to adapt rapidly to such changes. Existing technologies also tend to focus exclusively on missions that are continuously or repeatedly executed, but neglect short-term missions that are planned quickly and executed only once. The key challenge to dealing with both these deficiencies is that the network mapping process must be agile enough to determine cyber key terrain in a few hours. A focus on capability mapping could be a very effective mechanism to accelerate the mapping of short-term missions. We conclude that the ability to map network requirements of individual tasks or network capabilities is currently the most important technology gap.

We have proposed a number of technology thrusts that differ from the technologies we encountered in our survey. These focus on active or constructive approaches to cyber network mission dependency mapping, and provide potential avenues to map network tasks and capabilities to a greater level of detail. Investment in these or other technologies that merge manual and data-driven processes, and focus on connecting all of the information abstraction layers in Figure 2 (Missions, Processes, Capabilities and Logs/Artifacts) will undoubtedly improve the United States' ability to fight through a contested cyber environment, assess the mission impact of degraded or disabled assets, and provide superior mission assurance.

This page intentionally left blank.

GLOSSARY

MIT LL	MIT Lincoln Laboratory
MIT	Massachusetts Institute of Technology
AMMO	Automated Mission Mapping Operations
AOC	Air Operations Center
CAMUS	Cyber Assets to Missions and Users
CKT	Cyber Key Terrain
CMIA	Cyber Mission Impact Assessment
CNMD	Cyber Network Mission Dependencies
CPT	Cyber Protection Team
DoD	Department of Defense
FFRDC	Federally Funded Research and Development Center
GUI	Graphical User Interface
IT	Information Technology
JHU/APL	Johns Hopkins University Applied Physics Laboratory
JMPS	Joint Mission Planning System
JOPES	Joint Operation Planning and Execution System
LLNL	Lawrence Livermore National Laboratory
NetSPA	Network Security Planning Architecture
NeMS	Network Mapping System
NSDMiner	Mining for Network Service Dependencies
PCAMM	Person Centric Automated Mission Mapping
RiskMAP	Risk-to-Mission Assessment Process
SAB	Science Advisory Board
SME	Subject Matter Expert
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VLAN	Virtual Local Area Network
VM	Virtual Machine
YMAL	You May Also Like

This page intentionally left blank.

REFERENCES

[1] R. Gates, *Quadrennial Defense Review Report*, DIANE Publishing Company (February 2010), https://books.google.com/books?id=U4LUaIlbOoEC.

[2] C.A. Aitken, "Beyond Transformation: The CPO1/CWO Strategic Employment Model," produced for the Chief of Force Development by 17 Wing Winnipeg Publishing Office WPO30734 (2011).

[3] R. Elder, "Defending and Operating in a Contested Cyber Domain," Air Force Scientific Advisory Board, Winter Plenary (2008).

[4] K. Ingols, R. Lippmann, and K. Piwowarski, "Practical Attack Graph Generation for Network Defense," in *22nd Annual Computer Security Applications Conference (ACSAC'06)* (2006), 121–130.

[5] K.M. Carter, N. Idika, and W.W. Streilein, "Probabilistic Threat Propagation for Malicious Activity Detection," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2940–2944.

[6] A.E. Schulz, D. O'Gwynn, and J. Kepner, "Dynamically Correlating Network Terrain to Organizational Missions," in prep. (2014).

[7] S. Musman, M. Tanner, A. Temin, E. Elsaesser, and L. Loren, "A Systems Engineering Approach for Crown Jewels Estimation and Mission Assurance Decision Making," in *2011 IEEE Symposium on Computational Intelligence in Cyber Security (CICS)*, 210–216.

[8] J. Watters, S. Morrissey, D. Bodeau, and S.C. Powers, "The Risk-to-Mission Assessment Process (RiskMAP): A Sensitivity Analysis and an Extension to Treat Confidentiality Issues," in *Institute for Information Infrastructure* (July 2009).

[9] E. Peterson, "Dagger: Modeling and Visualization for Mission Impact Situation Awareness," manuscript submitted to *2015 IEEE International Symposium on Technologies for Homeland Security*.

[10] Lawrence Livermore National Laboratory, "NeMS: Network Mapping System," https://ipo.llnl.gov/technologies/nems.

[11] J.R. Goodall, A. D'Amico, and J.K. Kopylec, "Camus: Automatically Mapping Cyber Assets to Missions and Users," in *Military Communications Conference* (2009).

[12] A. Natarajan, P. Ning, Y. Liu, S. Jajodia, and S.E. Hitchinson, "NSDMiner: Automated Discovery of Network Service Dependencies," in *IEEE International Conference on Computer Communications* (March 2012).

[13] Lumeta, "Ipsonar," http://www.lumeta.com/product/ipsonar.html.

[14] A. Brown, G. Kar, and A. Keller, "An Active Approach to Characterizing Dynamic Dependencies for Problem Determination in a Distributed Environment," in *2001 IEEE/IFIP International Symposium on Integrated Network Management*, 377–390.

[15] Y. Izrailevsky and A. Tseitlin, "The Netflix Simian Army," *The Netflix Tech Blog* (July 2011).

This page intentionally left blank.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)* 18 September 2015	**2. REPORT TYPE** Technical Report	**3. DATES COVERED** *(From - To)*

4. TITLE AND SUBTITLE

Cyber Network Mission Dependencies

5a. CONTRACT NUMBER
FA8721-05-C-0002

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Alexia E. Schulz and Michael C. Kotson, Group 51
Joseph R. Zipkin, Group 58

5d. PROJECT NUMBER
2467

5e. TASK NUMBER
273

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

8. PERFORMING ORGANIZATION REPORT NUMBER

TR-1189

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Jim Bertone
Air Force Life Cycle Management Center, Development Planning,
Cyber & Space (AFLCMC/XZCC)
20 Schilling Circle, Bldg. 1305, Hanscom AFB, MA 01730

10. SPONSOR/MONITOR'S ACRONYM(S)
AFLCMC/XZCC

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Distribution Statement A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Cyber assets are critical to mission success in every arena of the Department of Defense (DoD). Because all DoD missions depend on cyber infrastructure, failure to secure network assets and assure the capabilities they enable will pose a fundamental risk to any defense mission. The impact of a cyber attack is not well understood by warfighters or leadership. It is critical that the DoD develop better cognizance of Cyber Network Mission Dependencies. This report identifies the major drivers for mapping missions to network assets, introduces existing technologies in the mission-mapping landscape, and proposes directions for future development.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT Unclassified	**b. ABSTRACT** Unclassified	**c. THIS PAGE** Unclassified	Same as report	35	**19b. TELEPHONE NUMBER** *(include area code)*